

CLAIMS

Listing of claims

1. (Previously Presented) A computer-implemented method for group communication over a network of processors comprising:
 - determining an overlay spanning tree comprising an origin node and at least one receiving node;
 - determining a configuration of the overlay spanning tree having a maximum throughput among all possible configurations of the overlay spanning tree, wherein determining the configuration of the overlay spanning tree comprises
 - defining a target bandwidth less than a maximum link bandwidth of edges of the overlay spanning tree given a fully connected overlay distribution graph,
 - constructing a reduced overlay distribution graph by iteratively removing an edge from a current overlay distribution graph, beginning with the fully connected overlay distribution graph, the edge having a bandwidth less than or equal to the target bandwidth, increasing the target bandwidth upon determining that the configuration of the overlay spanning tree is constructible based on the current overlay distribution graph, and decreasing the target bandwidth upon determining that the configuration of the overlay spanning tree is not constructible based on the current overlay distribution graph, until the configuration of the overlay spanning tree has the maximum throughput with no edge having a bandwidth below the target bandwidth; and
 - controlling a source communication rate between the origin node and the at least one

receiving node to be less than or equal to a bottleneck rate of the configuration of the overlay spanning tree based on the reduced overlay distribution graph.

2. (Original) The computer-implemented method of claim 1, further comprising protecting data delivery by link error recovery.

3. (Original) The computer-implemented method of claim 2, wherein the overlay spanning tree comprises a plurality of nodes, wherein the data delivery is reliable such that each node receives the same data.

4. (Original) The computer-implemented method of claim 1, further comprising scaling the overlay spanning tree to an arbitrary group size.

5-6. (Canceled)

7. (Original) The computer-implemented method of claim 1, further comprising joining a new node to the spanning tree.

8. (Original) The computer-implemented method of claim 7, comprising joining the new

node to an existing node of the spanning tree upon determining that the existing node has a bandwidth of greater than or equal to an existing rate.

9. (Original) The computer-implemented method of claim 8, further comprising:

determining a triangular improvement upon determining that no existing node has a bandwidth greater than or equal to the existing rate;

joining the new node at an attachment point having a highest bandwidth among existing nodes of the spanning tree upon determining that the triangular improvement failed; and

redetermining the spanning tree upon determining bandwidth less than or equal to a minimum threshold.

10. (Original) The computer-implemented method of claim 1, further comprising

redetermining the spanning tree upon determining that an existing node has left the spanning tree.

11. (Original) The computer-implemented method of claim 10, further comprising:

determining orphaned child nodes of the existing node that has left the spanning tree; and

performing a join for each orphaned child node.

12. (Previously Presented) A program storage device readable by machine, tangibly embodying a program of instructions executable by the machine to perform method steps for group communication over a network of processors, the method steps comprising:

 determining an overlay spanning tree comprising an origin node and at least one receiving node;

 determining a configuration of the overlay spanning tree having a maximum throughput among all possible configurations of the overlay spanning tree, wherein determining the configuration of the overlay spanning tree comprises

 defining a target bandwidth less than a maximum link bandwidth of edges of ~~for~~ the overlay spanning tree given a fully connected overlay distribution graph,

 constructing a reduced overlay distribution graph by iteratively removing an edge from a current overlay distribution graph, beginning with the fully connected overlay distribution graph, the edge having a bandwidth less than or equal to the target bandwidth, increasing the target bandwidth upon determining that the configuration of the overlay spanning tree is constructible based on the current overlay distribution graph, and decreasing the target bandwidth upon determining that the configuration of the overlay spanning tree is not constructible based on the current overlay distribution graph, until the configuration of the overlay spanning tree has the maximum throughput with no edge having a bandwidth below the target bandwidth; and

controlling a source communication rate between the origin node and the at least one receiving node to be less than or equal to a bottleneck rate of the configuration of the overlay spanning tree based on the reduced overlay distribution graph.

13. (Original) The method of claim 12, further comprising protecting data delivery by link error recovery.

14. (Original) The method of claim 13, wherein the overlay spanning tree comprises a plurality of nodes, wherein the data delivery is reliable such that each node receives the same data.

15. (Original) The method of claim 12, further comprising scaling the overlay spanning tree to an arbitrary group size.

16-17. (Canceled)

18. (Original) The method of claim 12, further comprising joining a new node to the spanning tree.

19. (Original) The method of claim 18, comprising joining the new node to an existing node of the spanning tree upon determining that the existing node has a bandwidth of greater than or equal to an existing rate.

20. (Original) The method of claim 19, further comprising:
determining a triangular improvement upon determining that no existing node has a bandwidth greater than or equal to the existing rate;
joining the new node at an attachment point having a highest bandwidth among existing nodes of the spanning tree upon determining that the triangular improvement failed; and
redetermining the spanning tree upon determining bandwidth less than or equal to a minimum threshold.

21. (Original) The method of claim 12, further comprising redetermining the spanning tree upon determining that an existing node has left the spanning tree.

22. (Original) The method of claim 21, further comprising:
determining orphaned child nodes of the existing node that has left the spanning tree; and
performing a join for each orphaned child node.

23. (Canceled)